WHAT IS CLAIMED IS:

1. A method for switching packets in a network switch, the method comprising:

receiving data forming a packet, wherein the packet is to be routed to at least one destination output port of a plurality of output ports that are part of the network switch;

determining whether the destination output port has sufficient resources available to handle the data without causing an overflow;

routing the data to the destination output port if sufficient resources are available; and

if sufficient resources are not available:

storing the data to a random access memory;

waiting until the destination output port has sufficient resources available; and

transferring the data from the random access memory to the destination output port.

2. The method of claim 1, wherein said routing is started before said receiving is complete to implement cut-through routing.

3. The method of claim 1, wherein said transferring is started before said receiving and said storing are complete to implement early forward routing.

4. The method of claim 1, wherein said reading is started after said storing is complete to implement store and forward routing.

5. The method of claim 1, wherein said routing is performed without storing the data in an intervening random access memory.

6. The method as recited in claim 1, further comprising:

storing the data received in the destination output port in an output first in first out memory (FIFO); and

determining whether the output FIFO has a predetermined amount of available storage space before performing said storing, wherein said predetermined amount of available storage space is stored in a programmable register.

7. The method as recited in claim 6, wherein said predetermined amount of available storage is a function of the maximum packet size.

8. The method as recited in claim 1, further comprising creating and storing a packet descriptor for the packet, wherein the packet descriptor is stored in a linked list.

9. The method as recited in claim 1, wherein said routing the data is performed to all output ports simultaneously that are destinations for the packet and that have sufficient resources available.

10. The method as recited in claim 1, wherein said transferring the data from the random access memory to the destination output port is performed once for each output port that is a destination for the packet.

11. The method as recited in claim 1, wherein said switching the packet comprises encapsulating the packet in a Storage Over IP (SOIP) packet if the packet is a Fibre Channel packet and the output port is an Ethernet port.

12. The method as recited in claim 1, wherein said storing comprises allocating clusters to the data, wherein each cluster comprises one or more cells, wherein each cell comprises a number of bytes equal to the width of the random access memory's interface.

13. A method for switching packets in a network switch, the method comprising:

receiving data forming a packet, wherein the packet is to be routed to at least one corresponding output port of a plurality of output ports that are part of the network switch; and

switching the packet by:

(a) if the corresponding output port has resources available to handle the data, cut-through routing the data to the corresponding output port; and

(b) if the corresponding output port does not have resources to handle the data, storing the data to a memory, and, in response to detecting that the memory has resources to handle the data, forwarding the packet to the output port from the memory, wherein the forwarding begins before the storing is complete.

14. A network switch comprising:

a plurality of input ports configured to receive data forming one or more packets;

a plurality of output ports configured to convey the packets out of the switch;

a random access memory; and

data transport logic coupled between the input ports, the output ports, and the memory, wherein a first input port is configured to request cut-through routing from at least one destination output port in response to receiving data corresponding to a first packet that is a candidate for cut-through routing,

wherein the des1tination output port is configured to convey a signal granting cut-through to the first input port if the destination output port has sufficient resources available to handle the data corresponding to the first packet,

wherein, in response to receiving the grant cut-through signal, the input port is configured to route the data corresponding to the first packet to the destination output port via the data transport logic, and wherein in response to not receiving the grant cut-through signal, the input port is configured to store the data to the shared memory via the data transport logic, wherein the output port is configured to read the data corresponding to the first packet from the shared memory via the data transport logic in response to having resources available for the data corresponding to the first packet.

15. The network switch of claim 14, wherein each output port comprises a plurality of output queues, wherein, in response to not receiving the grant cut-through signal, the input port is configured to store a packet identifier on one of the destination port's output queues.

16. The network switch as recited in claim 14, wherein each input port comprises an input port FIFO, wherein each output port comprises an output FIFO, wherein the input FIFO is configured to store received data until the received data is routed to the output FIFO or the random access memory.

17. The network switch of claim 14, wherein the first input port is configured to refrain from requesting cut-through routing for packets larger than the destination output port's output FIFO.

18. The network switch of claim 14, wherein one or more of the input ports and one or more of the output ports operate at different data rates.

19. The network switch of claim 14, wherein the destination output port is configured to ensure that a sufficient amount of data from the packet has been stored into the shared memory before starting to read the data corresponding to the first packet from memory, thereby preventing an output under-run.

20. The network switch of claim 14, wherein the sufficient amount of data is programmable on a port-by-port basis.

21. The network switch of claim 14, wherein the sufficient amount of data is a function of the output port's line rate.

22. The network switch as recited in claim 14, wherein the input port is configured to store the data to the shared memory via the data transport logic in cells.

23. The network switch as recited in claim 14, wherein the input port is configured to allocate one or more clusters for each received packet, wherein each cluster comprises one or more cells, wherein each cell equals the size of the random access memory's interface.

24. The network switch as recited in claim 14, wherein each output port comprises an output FIFO, wherein each output port is configured to refrain from conveying the signal granting cut-through in response to the output port's output FIFO being more than a predetermined percentage full.

25. The network switch as recited in claim 14, wherein each output port comprises an output FIFO, wherein each output port is configured to refrain from conveying the signal granting cut-through in response to the output port's output FIFO having less than a predetermined number of bytes of available storage.

26. The network switch of claim 14, wherein each output port comprises an output FIFO, wherein the output port comprises a means for delaying the forwarding of the stored data to the output FIFO from the RAM until the RAM has received enough of the packet to ensure that the output port will not underflow.

27. The network switch of claim 14, wherein each output port comprises an output FIFO, wherein the input port comprises an input FIFO and a means for delaying the cut-through routing of a particular packet from the input FIFO to the corresponding output port's output FIFO until the input port has received enough of the particular packet to ensure that the input FIFO will not underflow.

28. The network switch as recited in claim 14, wherein the data transport logic includes a cross-bar switch configurable to route data from each input port to each output port.

29. The network switch as recited in claim 14, wherein the crossbar switch is configured to prevent data from a particular input port being routed back to the particular input port.

30. The network switch as recited in claim 14, wherein said network switch further comprises a network processor for each input port and each output port, wherein said network processors are configured to add an Ethernet prefix to packets in response to detecting that the packets are Fibre Channel packets and are being routed to Ethernet output ports.

31. The network switch as recited in claim 14, wherein the packets have variable lengths.

32. The network switch as recited in claim 14, wherein said input port is configured to generate a packet descriptor for each received packet, wherein the packet descriptor comprises at least the length of the corresponding packet and an identifier that identifies the corresponding output port to which the packet is to be routed.

33. The network switch as recited in claim 14, wherein said input port is configured to generate a packet descriptor for the packet, wherein the packet descriptor comprises an indication of the packet's priority.

34. The network switch as recited in claim 14, wherein said switch is configurable to disable cut-through routing on a port-by-port basis.

35. The network switch as recited in claim 14, wherein said switch is configurable to disable cut-through routing on a packet-by-packet basis.

36. The network switch as recited in claim 14, wherein said switch is configurable to disable early forwarding on a port-by-port basis.

37. The network switch as recited in claim 14, wherein said switch is configurable to disable early forwarding operations on a packet-by-packet basis.

38. The network switch as recited in claim 14, further comprising a management CPU, wherein each of the plurality of input ports is configured to refrain from requesting cut-through for packets destined for the management CPU.

39. The network switch as recited in claim 14, wherein said network switch is configurable to selectively disable cut-through operations from a particular type of input port to a particular type of output port.

40. The network switch as recited in claim 14, wherein the input ports are either Fibre Channel or Gigabit Ethernet, and wherein the output ports are either Fibre Channel or Gigabit Ethernet.

41. A network switch comprising:

a plurality of ports, wherein a first one of said ports is an input port configured to receive data forming a packet, wherein a second one of said ports is an output port configured to convey the packet out of the switch, wherein the output port comprises an output first-in first-out memory (FIFO);

a random access memory (RAM); and

a means for routing the data between the input port, the RAM, and the output port, wherein the means for routing is configured to either route the packet directly to the output FIFO by cut-through routing or route the packet to the RAM for either early forwarding or store and forward routing.

42. The network switch of claim 41, wherein the routing means is configured to determine whether the output FIFO has sufficient storage available to store the packet, wherein the means is configured to route the data from the input port to the output port in response to detecting that the output FIFO has resources available for the packet, wherein the means is configured to store the data to the RAM in response to detecting that the output FIFO does not have room available for the packet, and wherein, in response to detecting that the output FIFO has room available to store the packet after at least a portion of the data has already been stored in the RAM, the means is configured to forward the stored data to the output FIFO.

43. The network switch of claim 42, wherein the forwarding is started before the entire packet is stored in the RAM.

44. The network switch of claim 42, wherein the forwarding is started before the entire packet is received by the input port.

45. The network switch as recited in claim 42, wherein the input port operates a different data rate than the output port.

46. The network switch of claim 42, wherein the output port comprises a plurality of output queues, wherein the input port is configured to store a packet identifier corresponding to the packet on one of the output port's output queues.

47. The network switch of claim 42, wherein the output port comprises a means for delaying the forwarding of the stored data to the output FIFO until the RAM has received enough of the packet to ensure that the output port will not underflow.

48. The network switch of claim 42, wherein the output port comprises control logic and an availability register, wherein the availability register is configured to store a value that the control logic compares with the amount of storage available in the output FIFO to determine if a signal granting cut-through should be transmitted to the input port.

49. The network switch of claim 48, wherein the value stored in the availability register is selected to prevent under-run on the output port's output FIFO when packets from a slower speed input port are routed to a faster speed output port.

50. The network switch of claim 42, wherein the input port is configured to refrain from requesting cut-through routing if the packet is larger than the output port's output FIFO.

51. The network switch as recited in claim 42, wherein the input port is configured to store the packet to the RAM by allocating clusters and cells to the packet.

52. The network switch as recited in claim 42, wherein said switch is configurable to selectively disable cut-through routing and early forwarding.

53. The network switch as recited in claim 42, further comprising a management CPU, wherein the input port is configured to refrain from requesting cut-through for packets destined for the management CPU.

54. The network switch as recited in claim 42, wherein the input port is either Fibre Channel or Gigabit Ethernet, and wherein the output port is either Fibre Channel and Gigabit Ethernet.